arXiv:1702.03375v1 [cs.DS] 11 Feb 2017

Derandomized Balanced Allocation

Xue Chen*
Computer Science Department
University of Texas at Austin
xchen@cs.utexas.edu

February 14, 2017

Abstract

In this paper, we study the number of random bits in the explicit constructions of hash functions and their maximum loads in the d-choice schemes when allocate sequentially n balls into n bins. We consider the *Uniform-Greedy* scheme [ABKU99], which provides d independent bins for each ball and places the ball into the bin with the least load, and its non-uniform variant — the *Always-Go-Left* scheme introduced by Vöcking [VÖ3]. We construct a hash function based on the previous work of Celis et al. [CRSW13] and show the following results.

1. This hash function with $O(\log n \log\log n)$ random bits has a maximum load of $\frac{\log\log n}{\log d} + O(1)$ with high probability in the *Uniform-Greedy* scheme, which matches the maximum load of a perfectly random hash function [ABKU99, VÖ3].
2. This hash function with $O(\log n \log\log n)$ random bits has a maximum load of $\frac{\log\log n}{d\log\phi_d} + O(1)$ with high probability in the *Always-Go-Left* scheme for a constant $\phi_d > 1.61$, which matches the maximum load of a perfectly random hash function [VÖ3].

Previously, the best known hash function that guarantees the same maximum loads as a perfectly random hash function in these two schemes was $O(\log n)$-wise independent functions [VÖ3], which needs $\Theta(\log^2 n)$ random bits.

*Supported by NSF Grant CCF-1526952 and a Simons Investigator Award (#409864, David Zuckerman).

1 Introduction

We investigate explicit constructions of hash functions for the classical problem of placing balls into bins. The basic model is to hash n balls into n bins independently and uniformly at random, which we call 1-choice scheme. A well-known and useful fact of the 1-choice scheme is that with high probability, each bin contains at most $O(\log n/\log\log n)$ balls. For convenience, we always use logarithm of base 2. Here, by high probability, we mean probability $1-n^{-c}$ for an arbitrary constant c.

An alternative variant, which we call *Uniform-Greedy*, is to provide $d \geq 2$ independent random choices for each ball and place the ball in the bin with the lowest load. In a seminal work, Azar et al. [ABKU99] showed that the *Uniform-Greedy* scheme with d independent random choices guarantees a maximum load of only $\frac{\log\log n}{\log d} + O(1)$ with high probability for n balls. Later, Vöcking [VÖ03] introduced the *Always-Go-Left* scheme to further improve the maximum load to $\frac{\log\log n}{d\log\phi_d} + O(1)$ for d choices where $\phi_d > 1.61$ is the constant satisfying $\phi_d^d = 1 + \phi_d + \cdots + \phi_d^{d-1}$. For convenience, we always use d-choice schemes to denote the *Uniform-Greedy* scheme and the *Always-Go-Left* scheme with $d \geq 2$ choices.

Traditional analysis of load balancing assumes a perfectly random hash function. A large body of work is dedicated to the removal of this assumption by designing explicit hash functions using fewer random bits. In the 1-choice scheme, it is well known that $O(\frac{\log n}{\log\log n})$-wise independent functions guarantee a maximum load of $O(\frac{\log n}{\log\log n})$ with high probability, which reduces the number of random bits to $O(\frac{\log^2 n}{\log\log n})$. Recently, Celis et al. [CRSW13] designed a hash function with a description of $O(\log n\log\log n)$ random bits that achieves the same maximum load of $O(\frac{\log n}{\log\log n})$ as a perfectly random hash function.

In this work, we are interested in the explicit constructions of hash functions with a few random bits that achieve *the same maximum loads* as a perfectly random hash function in the d-choice schemes. More precisely, we study how to derandomize the perfectly random hash function whose maximum loads are $\frac{\log\log n}{\log d} + O(1)$ in the *Uniform-Greedy* scheme [ABKU99, VÖ03] and $\frac{\log\log n}{d\log\phi_d} + O(1)$ in the *Always-Go-Left* scheme [VÖ03]. For these two d-choice schemes, $O(\log n)$-wise independent hash functions suffice to achieve the same maximum loads with high probability from Vöcking's argument [VÖ03], which provides a hash function with $\Theta(\log^2 n)$ random bits. Very recently, Reingold et al. [RRW14] showed that the hash function by Celis et al. [CRSW13] guarantees a maximum load of $O(\log\log n)$ in the *Uniform-Greedy* scheme with $d \geq 2$ choices using $O(\log n\log\log n)$ random bits.

1.1 Our Contributions

We first show a hash function with an explicit description of $O(\log n\log\log n)$ random bits that guarantees a maximum load of $\frac{\log\log n}{\log d} + O(1)$ in the *Uniform-Greedy* scheme [ABKU99, VÖ03] with d choices. We use U to denote the pool of balls and consider placing $m = O(n)$ balls into n bins here. Without loss of generality, we always assume $|U| = \mathsf{poly}(n)$ and $d \geq 2$ in this work.

Theorem 1.1 (Informal version of Theorem 5.1) *For any $m = O(n)$, any constants c and d, there exists a hash function h with $O(\log n\log\log n)$ random bits such that given any m balls in U, with probability at least $1-n^{-c}$, the max-load of the* Uniform-Greedy *scheme with d independent choices of h is $\frac{\log\log n}{\log d}+O(1)$.*

We prove our main theorem by derandomizing Vöcking's argument [VÖ03] for the d-choice schemes carefully. We introduce new ideas to strengthen the hash function designed by Celis et al. [CRSW13] for the 1-choice scheme such that we could adopt Vöcking's counting argument [VÖ03]. Our hash function has an evaluation time $O\big((\log\log n)^4\big)$ in the RAM model based on the algorithm designed by Meka et al. [MRRR14].

Then we show that this hash function guarantees a maximum load of $\frac{\log\log n}{d\log\phi_d}+O(1)$ in the *Always-Go-Left* scheme [VÖ03] with d choices. The *Always-Go-Left* scheme [VÖ03] is an asymmetric allocation scheme that partitions the n bins into d groups with equal size and uses an unfair tie-breaking mechanism. The allocation process of the *Always-Go-Left* scheme provides d independent choices for each ball from the d groups separately and always chooses the left-most bin with the least load for each ball. We defer the formal description of the *Always-Go-Left* scheme to Section 6. Observe that the constant ϕ_d in equation $\phi_d^d = 1+\phi_d+\cdots+\phi_d^{d-1}$ satisfies $1.61 < \phi_2 < \phi_3 < \phi_4 < \cdots < \phi_d < 2$. This indicates that the *Uniform-Greedy* scheme improves the maximum load by a constant factor. For example, when $d=2$, the maximum load of the *Always-Go-Left* scheme is at most $0.7\log\log n+O(1)$.

Theorem 1.2 (Informal version of Theorem 6.3) *For any $m=O(n)$, any constants c and d, there exists a hash function h with $O(\log n\log\log n)$ random bits such that given any m balls in U, with probability at least $1-n^{-c}$, the max-load of the* Always-Go-Left *scheme with d independent choices of h is $\frac{\log\log n}{d\log\phi_d}+O(1)$.*

Finally, we show that our hash function guarantees the same maximum load in the 1-choice scheme as a perfectly random hash function for $m=n\cdot\mathsf{poly}(\log n)$ balls. It is well known that given $m>n\log n$ balls in U, the maximum load is $\frac{m}{n}+O(\sqrt{\log n\cdot\frac{m}{n}})$ for the 1-choice scheme from the Chernoff bound. For convenience, we refer to this case of $m\geq n\log n$ balls as a *heavy load*. Very recently, Gopalan, Kane, and Meka [GKM15] designed a pseudorandom generator of seed length $O(\log n(\log\log n)^2)$ that fools the Chernoff bound within polynomial error. Hence the pseudorandom generator of [GKM15] provides a hash function with $O(\log n(\log\log n)^2)$ random bits achieving a maximum load of $\frac{m}{n}+O(\sqrt{\log n\cdot\frac{m}{n}})$ in the heavy load case. Compared to the hash function of [GKM15], we provide a simplified construction that achieves this maximum load but only works for $m=n\cdot\mathsf{poly}(\log n)$ balls.

Theorem 1.3 (Informal version of Theorem 7.1) *For any constants c and $a\geq 1$, there exist a hash function generated by $O(\log n\log\log n)$ random bits such that for any $m=\log^a n\cdot n$ balls, with probability at least $1-n^{-c}$, the max-load of the n bins in the 1-choice scheme with h is $\frac{m}{n}+O\left(\sqrt{\log n}\cdot\sqrt{\frac{m}{n}}\right)$.*

1.2 Previous Work

scheme	reference	maximum load	number of random bits
1-choice	well known	$O(\frac{\log n}{\log\log n})$	$\Theta(\frac{\log^2 n}{\log\log n})$
1-choice	[CRSW13]	$O(\frac{\log n}{\log\log n})$	$O(\log n\log\log n)$
Uniform-Greedy	[ABKU99]	$\frac{\log\log n}{\log d}+O(1)$	full randomness
Uniform-Greedy	[VÖ03]	$\frac{\log\log n}{\log d}+O(1)$	$\Theta(\log^2 n)$
Uniform-Greedy	[RRW14]	$O(\log\log n)$	$O(\log n\log\log n)$
Uniform-Greedy	[DKRT16]	$O(\log\log n)$	$\Theta(\log^2 n)$
Uniform-Greedy	this work	$\frac{\log\log n}{\log d}+O(1)$	$O(\log n\log\log n)$
Always-Go-Left	[VÖ03]	$\frac{\log\log n}{d\log\phi_d}+O(1)$	$\Theta(\log^2 n)$
Always-Go-Left	[Woe06]	$\frac{\log\log n}{d\log\phi_d}+O(1)$	$n^{\Theta(1)}$
Always-Go-Left	this work	$\frac{\log\log n}{d\log\phi_d}+O(1)$	$O(\log n\log\log n)$

Table 1: Summary of previous works about the maximum loads of placing n balls in n bins

Natural explicit constructions of hash functions using a few random bits are k-wise independent functions, small-biased spaces, and k-wise small-biased spaces. For the 1-choice scheme with $m=n$ balls,

Alon et al. [ADM+99] showed the existence of a pairwise independent hash family that always has a maximum load of $\sqrt{n}$. On the other hand, it is well known that $O(\frac{\log n}{\log \log n})$-wise independent functions achieve a maximum load of $O(\frac{\log n}{\log \log n})$ with high probability, which needs $O(\frac{\log^2 n}{\log \log n})$ random bits. Using $O(\log n)$-wise small-biased spaces as milder restrictions, Celis et al. [CRSW13] designed a hash function with $O(\log n \log \log n)$ random bits achieving the same maximum load with high probability.

For the heavy load case in the 1-choice scheme, the maximum load is $\frac{m}{n} + O(\sqrt{\log n \cdot \frac{m}{n}})$ from the Chernoff bound. Hence any pseudorandom generator fooling the Chernoff bound within polynomial small error is a hash function matching this maximum load. Schmidt et al. [SSS95] showed that $O(\log n)$-wise independent functions could derandomize the Chernoff bound, which provides a hash function with $O(\log^2 n)$ random bits. In a recent breakthrough [GKM15], Gopolan, Kane, and Meka designed a pseudorandom generator with seed length $O(\log n (\log \log n)^2)$ to fool halfspaces, the Chernoff bound, and many other classes, which provides a hash function using $O(\log n (\log \log n)^2)$ random bits.

For $m = n$ balls in the d-choice schemes, the original argument of [ABKU99] adopts an inductive proof that relies on the assumption of full randomness. It is known (e.g., [RRW14, DKRT16]) that $O(\log n)$-wise independent functions could derandomize Vöcking's counting argument [VÖ03], which takes $\Theta(\log^2 n)$ random bits to achieve the maximum load of $\frac{\log \log n}{\log d} + O(1)$ in the *Uniform-Greedy* scheme and $\frac{\log \log n}{d \log \phi_d} + O(1)$ in the *Always-Go-Left* scheme. Very recently, Reingold et al. [RRW14] proved that the hash function in [CRSW13] achieves a maximum load of $C \cdot \log \log n + O(1)$ in the *Uniform-Greedy* scheme with 2 choices, where $C > 2$ is the universal constant such that in a random graph with n vertices and n/C edges, with high probability, every connected component has size $O(\log n)$.

Another research direction of hash functions focuses on studying functions with a constant evaluation time despite the expense of the number of random bits. Woelfel [Woe06] showed that the hash function of [DW03], which takes constant evaluation time and $n^{\Theta(1)}$ random bits, guarantees the same maximum loads as a perfectly random hash functions in the two d-choices schemes. Pătraşcu and Thorup [PT12] introduced simple tabulation hashing, a hash function with constant evaluation time and $n^{\Theta(1)}$ random bits that can replace the perfectly random hash functions in various applications. Very recently, Dahlgaard [DKRT16] proved that the maximum load of the *Uniform-Greedy* scheme is $O(\log \log n)$ with high probability in simple tabulation [PT12] using $O(\log^2 n)$ random bits. For the hash function in [CRSW13], Meka et al. [MRRR14] improved its evaluation time to $O((\log \log n)^2)$ in the RAM mode.

We summarize these results in Table 1. Finally, we refer surveys [Mit01, MRS00] and the reference therein for various applications of multiple-choice schemes in computer science.

1.3 Organization

This paper is organized as follows. We provide a brief overview about our hash function and sketch the proof of the maximum load in the *Uniform-Greedy* scheme in Section 2. In Section 3, we introduce some notations and tools then review Vöcking's argument [VÖ03]. We show the construction of our hash functions in Section 4. Next we prove Theorem 1.1 in Section 5 and Theorem 1.2 in Section 6, which provide upper bounds on the maximum load of the *Uniform-Greedy* scheme and the maximum load of the *Always-Go-Left* scheme separately. Finally, we prove Theorem 1.3 in Section 7 which shows a bound of the heavy load case in the 1-choice scheme.

2 Proof Overview

We sketch our proofs in this section. We first define our hash function and point out the main properties that will be used in this work. Then we outline the proof for the *Uniform-Greedy* scheme. Without loss of generality, we assume n is a power of 2 and fix $m = n$ balls in this section. For convenience, we always use $[n]$ to denote the subset $\{1, 2, \cdots, n\}$.

Our construction: Our hash function is an extension of the hash function in [CRSW13]. We define our hash function $h : U \to [n]$ to be

$$h(x) = \big(h_1(x) \circ h_2(x) \circ \cdots \circ h_k(x) \circ h_{k+1}(x)\big) \oplus g(x)$$

where

1. $\circ$ denotes concatenation. $\oplus$ stands for a bit-wise XOR operation,
2. $k = \log(\log n/3 \log\log n)$ such that $n^{2^{-k}} = \log^3 n$,
3. h_i is a δ_1-biased space from U to $[n^{2^{-i}}]$ for each $i \in [k]$ with $\delta_1 = 1/\mathsf{poly}(n)$,
4. h_{k+1} is a δ_2-biased space from U to $[n^{2^{-k}}] = [\log^3 n]$ with $\delta_2 = (\log n)^{-C \log n}$ for a constant C,
5. $g(x)$ is k_g-wise independent from U to $[n]$ for $k_g = \Theta(\log\log n)$.

Thus our hash function takes $O(k \log \frac{|U| \log n}{\delta_1} + \log \frac{|U| \log n}{\delta_2} + k_g \log n) = O(\log n \log\log n)$ random bits.

We state several key properties of h. Because g is $\Theta(\log\log n)$-wise independent, we have the same property for h.

Claim 2.1 *h is $O(\log\log n)$-wise independent.*

Notice that the subfunctions $h_1, h_2, \cdots, h_{k+1}$ constitute a function from U to $[n]$ because the images satisfy $\prod_{j=1}^{k} [n^{2^{-j}}] \times [n^{2^{-k}}] = [n]$. For $h_1 \circ h_2 \cdots \circ h_{k+1}$, we consider it as an allocation process (after fixing g) that first allocates the n balls into $n^{1/2}$ bins by h_1. The next hash table h_2 further allocates the balls in every bin of $[n^{1/2}]$ (the image of h_1) to $n^{1/4}$ bins. In another word, $h_1 \circ h_2$ allocates n balls into $n^{1-\frac{1}{4}}$ bins, and so on for $h_1 \circ h_2 \cdots \circ h_i$.

After fixing g, $h_1 \circ h_2 \circ \cdots \circ h_k$ allocates the n balls into $[n^{1-2^{-k}}] = [n/\log^3 n]$ bins. Hence each bin contains $\log^3 n$ balls in expectation. Celis et al. [CRSW13] shows that the number of balls in each bin of $h_1 \circ h_2 \circ \cdots \circ h_k$ is very close to the expecation in the 1-choice scheme, which is stated as follows.

Claim 2.2 (Celis et al. [CRSW13]) *With high probability over $h_1 \circ h_2 \circ \cdots \circ h_k$, there are at most $1.01 \cdot \log^3 n$ balls in every bin of $[n/\log^3 n]$.*

Lemma 4.3 and Corollary 4.4 in Section 4 contain the formal statements and full proofs.

In the last subfunction, $h_{k+1} : U \to [\log^3 n]$ is a $\delta_2 = (\log n)^{-C \cdot \log n}$-biased space for a sufficiently large constant C. Before stating the main property of h_{k+1}, we introduce several notations. Let $\|X\|_1$ denote the ℓ_1-norm such that $\|X - Y\|_1 = \sum_{x \in \mathsf{supp}(X)} |X(x) - Y(x)|$ for a pair of random variables (X, Y) with the same support. For a subset $S \subseteq U$, we abuse the notation $D_{[\log^3 n]^S}$ denote the uniform distribution over all maps from $S \to [\log^3 n]$ here. From the property of small-biased spaces proved by Vazirani [Vaz86] and the union bound, we have the following claim.

Claim 2.3 *For a fixed subset S of size $\frac{C}{3}\cdot\log n$, $h_{k+1}(S)$ is $(\log n)^{-\frac{C}{2}\cdot\log n}$-close to the uniform distribution on S, i.e., $\|h_{k+1}(S) - D_{[\log^3 n]^S}\|_1 \le (\log n)^{-\frac{C}{2}\cdot\log n}$.*

For $m = (\log n)^{\frac{C}{3}\cdot\log n}$ subsets $S_1,\cdots,S_m$ of size $\frac{C}{3}\cdot\log n$, we have

$$\sum_{i\in[m]} \|h_{k+1}(S_i) - D_{[\log^3 n]^{S_i}}\|_1 \le m\cdot(\log n)^{-\frac{C}{2}\cdot\log n} \le (\log n)^{-\frac{C}{6}\cdot\log n}.$$

Notice that h_{k+1} (or h) is not close to $\log n$-wise independent functions on those n balls because there are $\binom{n}{\log n}$ subsets of size $\log n$.

Now we define the *Uniform-Greedy* scheme and consider its maximum load.

Definition 2.4 (*Uniform-Greedy* with d choices) *The allocation process works as follows: Let $h^{(1)},\cdots,h^{(d)}$ be d hash functions from U to $[n]$. For each ball i, the algorithm considers d bins $\{h^{(1)}(i),\cdots,h^{(d)}(i)\}$ and puts the ball i into the bin with the least load among $\{h^{(1)}(i),\cdots,h^{(d)}(i)\}$. When there are several bins with the least load, we pick an arbitrary one.*

For convenience, we define the height of a ball to be the height of this ball on the bin allocated in the above process. Vöcking [VÖ03] defined a witness tree for every ball b in the allocation process:

1. A ball b with height $l+4$ $\big($for some $l = \log_d\log n + O(1)\big)$ corresponds to a d-ary witness tree T of height l whose leaves have height at least 4.
2. Each node in the witness tree corresponds to a ball where the root corresponds to b.
3. A ball v is the ith child (for $i\in[d]$) of a ball u iff the ball v is the top ball of the bin $h^{(i)}(u)$ when we place u. This happens only if

$$h^{(i)}(u) \in \{h^{(1)}(v), h^{(2)}(v),\cdots,h^{(d)}(v)\}. \tag{1}$$

Given a tree T, we always use $|T|$ to denote the number of vertices in T for convenience.

Now we trim the repeated nodes in a witness trees and define *collisions* for repeated nodes. Given a witness tree T where nodes $v_1,\cdots,v_j$ in T correspond to the same ball, we always keep the first node v_1 in the breadth-first search of T. For the other nodes $v_2,\cdots,v_j$, we redirect the edges from their parents to v_1 and call these edges *collisions*. Then we remove $v_2,\cdots,v_j$ and their subtrees.

For convenience, we call different witness trees with the same structure but different balls as a *configuration*. Notice that the configuration of witness trees with distinct nodes is a full d-ary tree without any collision.

To illustrate our main ideas in the derandomization, we consider two cases: witness trees with distinct balls and witness trees with at least c collisions (for some constant $c>1$). We first outline the argument for distinct balls then outline the proof for at least c collisions.

Witness trees with distinct balls: We first revisit Vöcking's argument for witness trees with distinct balls in the uniform distribution. Recall that a ball with height $l+4$ corresponds to a witness tree with height l whose leaves have height at least 4. We bound the probability that a leaf ball has height at least 4 by 3^{-d} as follows. A leaf ball of height at least 4 indicates that each bin in his choices has height at least 3. Because at most $n/3$ bins contain at least 3 balls at any moment, the probability that a random bin has height at least 3 is $\le 1/3$. Thus the probability that d random bins have height 3 is at most 3^{-d}.

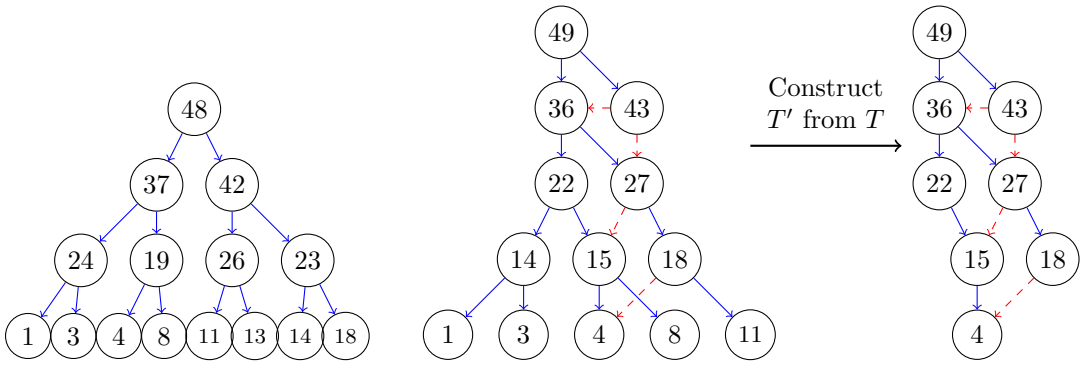

Figure 1: A witness tree with distinct balls and a witness tree with 4 collisions for 2 choices

We bound the probability that any witness tree of height l with distinct nodes exists in d perfectly random hash functions by

$$n^{|T|} \cdot \Big(\prod_{(u,v)\in T} \frac{d}{n}\Big) \cdot (3^{-d})^{\text{number of leaves}} \leq n \cdot d^{|T|-1} \cdot 3^{-d\cdot d^l} \leq n \cdot (d^2 3^{-d})^{d^l}, \tag{2}$$

where $n^{|T|}$ comes from the number of possible choices of balls for each node, and $(\frac{d}{n})$ is the probability of equation (1) for every edge. Hence the probability is polynomially small when $d^l = \Theta(\log n)$, which indicates that there is no ball of height $l + 4 = \log_d \log n + O(1)$.

Now we consider derandomzing the above argument. Let $h^{(1)}, \cdots, h^{(d)}$ be d independent hash functions in our construction. There are two probabilities in the witness trees argument (2): the second term $\prod_{(u,v)\in T} \frac{d}{n}$ about all edges in T and the last term $3^{-d\cdot d^l}$ about all leaves. The second term needs $O(\log n)$-wise independence over $[n]$ bins, which we cannot support with $o(\log^2 n)$ bits.

Our strategy is to fix the prefixes in the hash functions — $(h_1^{(1)} \circ h_2^{(1)} \cdots \circ h_k^{(1)}), (h_1^{(2)} \circ h_2^{(2)} \cdots \circ h_k^{(2)}), \cdots$ and $(h_1^{(d)} \circ h_2^{(d)} \cdots \circ h_k^{(d)})$ but leave the suffixes $h_{k+1}^{(1)}, \cdots, h_{k+1}^{(d)}$ unfixed (we do not need $g^{(1)}, \cdots, g^{(d)}$ in this case). We assume that the property in Claim 2.2 holds for all hash functions $h^{(1)}, \cdots, h^{(d)}$ in the rest of this discussion.

Notice that for an edge (u, v) in the witness tree T to satisfy (1), the prefixes of $h^{(1)}(v), \cdots, h^{(d)}(v)$ and $h^{(i)}(u)$ must satisfy

$$h_1^{(i)}(u) \circ \cdots \circ h_k^{(i)}(u) \in \left\{h_1^{(1)}(v) \circ \cdots \circ h_k^{(1)}(v), \cdots, h_1^{(d)}(v) \circ \cdots \circ h_k^{(d)}(v)\right\}. \tag{3}$$

Let $\mathcal{T}$ denote the subset of possible witness trees of height l that satisfy the above condition (3) for every edge after fixing the prefixes. Because there are at most $1.01 \log^3 n$ elements in each bin of $h_1^{(j)} \circ h_2^{(j)} \cdots \circ h_k^{(j)}$ for every function $h^{(j)}$ from Claim 2.2, we bound

$$|\mathcal{T}| \leq n(d \cdot 1.01 \log^3 n)^{|T|-1} = n \cdot (1.01d)^{|T|} \cdot (\log^3 n)^{|T|-1}$$

instead of $n^{|T|}$ in the original argument.

Next, the suffixes $h_{k+1}^{(1)}, \cdots, h_{k+1}^{(d)}$ also need to satisfy (1). We assume $O(\log n)$-wise independence in h_{k+1} and calculate the probability that these suffixes satisfy (1) for every edge in a fixed possible witness

tree among $\mathcal{T}$ (namely the second term in (2)) as $\prod_{(u,v)\in T} \frac{d}{\log^3 n} = (\frac{d}{\log^3 n})^{|T|-1}$. Thus the product of the first term and the second term in (2) becomes

$$|\mathcal{T}| \cdot (\frac{d}{\log^3 n})^{|T|-1} = n \cdot \mathsf{poly}(d)^{|T|}.$$

Similarly, we use this strategy to recalculate the last term 3^{-d} of every leaf ball in (2) by losing a constant on the height such that the union bound of the product becomes $n \cdot (\frac{1}{2})^{|T|} \leq 1/\mathsf{poly}(n)$ given $|T| = \Theta(\log n)$. Finally, we replace the $O(\log n)$-wise independence in the above calculation by a $(\log n)^{-C \log n}$-biased space. We use Claim 2.3 and the bound $|\mathcal{T}| = (\log n)^{O(\log n)}$ to finish the proof.

Witness trees with a lot of collisions: We fix a witness tree T and consider the first c collisions $e_1, \cdots, e_c$ in the BFS of T. Let T' be the subtree of T containing all nodes on $e_1, \cdots, e_c$ and their ancestors. Because the height of T is $l = O(\log \log n)$, the number of edges in T' is at most $c \cdot 2l + c = O(\log \log n)$.

T appears in the hash functions $h^{(1)}, h^{(2)}, \cdots, h^{(d)}$ only if T' appears in these hash functions. Because h is $k_g = \Theta(\log \log n)$-wise independent from Claim 2.1, we could bound the probability of any possible witness tress in the same configuration (structure) of T' appearing in h by:

$$n^{|T'|} \cdot (\frac{d}{n})^{|T'|-1+c}.$$

At the same time, the number of configurations (structures) with c collision is at most $(d^{l+1})^{2c}$. We apply a union bound on all possible configurations:

$$n^{|T'|} \cdot (\frac{d}{n})^{|T'|-1+c} \cdot (d^{l+1})^{2c} \leq n^{-c/3}.$$

Hence with probability $1 - n^{-c/3}$, there is no witness tree of height l and at least c collisions.

Generalizations: We will adopt this argument for the *Always-Go-Left* scheme. For the heavy load case of $m > n \log n$ balls, we choose the image set of h_{k+1} based on $\frac{m}{n}$ properly.

To improve the evaluation time of the hash function, we will replace the δ-biased spaces of the hash function h in this section by $\mathsf{poly}(\log n)$-wise δ-biased spaces such that we could evaluate h in $O((\log \log n)^4)$ operations in the RAM mode based on [MRRR14] without affecting out analysis.

3 Preliminaries

We use U to denote the pool of balls, m to denote the numbers of balls in U, and n to denote the number of bins. We assume $|U| = \mathsf{poly}(n)$ and $m \geq n$. Without loss of generality, we always assume d is a constant and n is a power of 2 in this work. We use 1_E to denote the indicator function of the event E.

Definition 3.1 *A distribution D on $\{0,1\}^n$ is a δ-biased space if for any non-trivial character function χ_S in $\{0,1\}^n$, $\mathrm{E}_{x\sim D}[\chi_S(x)] \leq \delta$.*

A distribution D on $\{0,1\}^n$ is a k-wise δ-biased space if for any non-trivial character function χ_S in $\{0,1\}^n$ of support size at most k, $\mathrm{E}_{x\sim D}[\chi_S(x)] \leq \delta$.

Vazirani [Vaz86] pointed out that small-biased spaces are close to the uniform distribution and the k-wise independent distributions. Given two distributions D_1 and D_2 on the same support $\{0,1\}^n$, we define the distance in ℓ_1 norm to be $\|D_1 - D_2\|_1 = \sum_{x\in\{0,1\}^n} |D_1(x) - D_2(x)|$. Given a distribution D_1 and a subset $\mathcal{D}$ of distributions, we define the distance between D_1 and $\mathcal{D}$ to be $\min_{D_2\in\mathcal{D}}\{\|D_1 - D_2\|_1\}$.

Lemma 3.2 *A δ-biased space on $\{0,1\}^n$ is $\delta \cdot 2^{n/2}$ close to the uniform distribution in ℓ_1 norm.*

Given a subset S of size k in $[n]$, a k-wise δ-biased space on $\{0,1\}^n$ is $\delta \cdot 2^{k/2}$ close to the uniform distribution on S in ℓ_1 norm.

The seminal works [NN90, AGHP90] provide pseudorandom generators of δ-biased spaces with seed length $O(\log \frac{n}{\delta})$ and k-wise δ-biased spaces with seed length $O(\log \frac{k \log n}{\delta})$.

In this work, for a distribution D on functions from U to $[n]$ where n is a power of 2, we say that D is k-wise independent if for any k elements $x_1, \cdots, x_k$ in U, $D(x_1, \cdots, x_k)$ is a uniform distribution on $[n]$. At the same time, we say that D is a k-wise δ-biased space if the distribution on the concatenated output of $|U| \cdot \log_2 n$ bits induced by D is a $(k \log_2 n)$-wise δ-biased space. From all discussion above, we have the following lemma.

Lemma 3.3 *Given k balls in U and n bins, a k-wise δ-biased space from U to $[n]$ is $\delta \cdot n^{k/2}$ close to the uniform distribution from U to $[n]$ on these k balls, which needs $O(\log \frac{k \log^2 n}{\delta})$ random bits.*

Given the domain U and the number of bins $[n]$, our hash functions are constituted by small-biased spaces and k-wise small-biased spaces in this work.

We state the Chernoff bound in k-wise independence by Schmidt et al. in [SSS95].

Lemma 3.4 (Theorem 5 (I) (b) in [SSS95]) *If X is the sum of k-wise independent random variables, each of which is confined to the interval $[0,1]$ with $\mu = \mathrm{E}[X]$, then for $\delta \le 1$ and $k \ge \delta^2 \mu \cdot e^{-1/3}$,*

$$\Pr[|X - \mu| \ge \delta\mu] \le e^{-\delta^2\mu/3}.$$

3.1 The Witness Tree argument [VÖ03]

We review the argument of Vöcking [VÖ03] for the *Uniform-Greedy* scheme.

Definition 3.5 (Symmetric Witness Tree) *A symmetric witness tree T with height l for a ball b is a complete d-ary tree of height l. Every node w in this tree corresponds to a ball $T(w) \in [n]$, and the root corresponds to the ball b. A ball u in T has a ball v as its ith child iff when we allocate u in the process, ball v is the top ball in the bin $h^{(i)}(u)$. Hence $v < u$ and the bin $h^{(i)}(u)$ is in the subset $\{h^{(1)}(v), \cdots, h^{(d)}(v)\}$ of $[n]$ when v is the ith child of u.*

Now we review the proof of [VÖ03] that with probability at least $1 - n^{-c}$, there is no ball of height $\log_d \log n + 3c + O(1)$ in a perfectly random hash function. For simplicity, we assume that there are $m = n$ balls in this calculation. We separate all witness trees into two cases: witness tree with at most $3c$ collisions and witness tree with at least $3c$ collisions.

Witness trees with at most $3c$ collision: Let us consider all witness trees of height $l = \log_d \log n + 3c$ whose leaves of height at least 4. We fix a configuration (structure of witness trees) C. Then each node of this configuration C corresponds a ball such that there are at most $n^{|C|}$ possible witness trees of height l. Next, we fix one possible witness tree T in the configuration C and consider all edges in the tree T. We

know that each edge (u, v) in T holds in the process if $h^{(i)}(u) \in \{h^{(1)}(v), \cdots, h^{(d)}(v)\}$, which happens with probability at most $\frac{d}{n}$.

Next, a leaf v has height at least 4 iff every choice from the d hash functions has height at least 3, which happens with prob. at most 3^{-d} from Section 2. We lower bound the number of leaves by $\frac{|C|-3c}{2}$ because C is a $(d \geq 2)$-ary tree except $3c$ nodes. Another lower bound of the number of leaves is $d^{l-3c} \geq (1+c)\log n$. We apply a union bound on the probability of a ball of height more than $l+4$:

$$n^{|C|} \cdot \prod_{(u,v)\in T} \frac{d}{n} \cdot 3^{-d \cdot d^l} \leq n \cdot d^{|C|} \cdot 3^{-d \cdot \frac{|C|-3c}{2}}.$$

For any constants $d \geq 2$ and $c > 2$, this probability is less than n^{-2c} when $|C| = \Omega(\log n)$. Finally, we bound the number of configurations by $\sum_{i=0}^{3c}(d^{l+1})^{2\cdot i} \leq n^{c/10}$. From the union bound, we bound the probability that any witness with height l and at most $3c$ collisions by n^{-c}.

Witness trees with at least $3c$ collision: We fix a configuration C with $3c$ collisions (we omit the rest collisions). Hence there are $|C|$ nodes in the configuration and $|C|-1+3c$ edges in the configuration. We apply a union bound on all possible choices of balls in C:

$$n^{|C|} \cdot (\frac{d}{n})^{|C|-1+3c} \leq n^{-1.5c}.$$

Then we count the number of configurations with $3c$ collisions, i.e., $(d^{l+1})^{2\cdot 3c} \leq n^{c/10}$. From the union bound, we bound the probability that any witness with height l and at least $3c$ collisions by n^{-c}.

$O(\log n)$-wise independence: At the same time, this argument also applies to independent hash functions $h^{(1)}, \cdots, h^{(d)}$ with $O(\log n)$-wise independence by losing a constant on the height.

4 The Construction of Our Hash Functions

We construct our hash function based on the hash function of [CRSW13] and provide the analysis for $m = O(n)$ balls in the 1-choice scheme to illustrate the key property that will be used in this work.

Definition 4.1 *Given $\delta_1 > 0, \delta_2 > 0$, and two integers k, k_g, we define $h : U \to [n]$ as follows:*

$$h(x) = \big(h_1(x) \circ h_2(x) \cdots \circ h_k(x) \circ h_{k+1}(x)\big) \oplus g(x).$$

where $\circ$ denotes concatenation and $\oplus$ denotes a bit-wise XOR. We specify each component in h:

1. $h_i : U \to [n^{2^{-i}}]$ is an $O(\log n)$-wise δ_1-biased space for each $i \in [k]$.
2. $h_{k+1} : U \to [n^{2^{-k}}]$ is an $O(\log n \log\log n)$-wise δ_2-biased space such that $(h_1(x) \circ h_2(x) \cdots \circ h_k(x) \circ h_{k+1}(x))$ is a function from U to $[n]$.
3. $g : U \to [n]$ is k_g-wise independent from U to $[n]$.

Hence the seed length of h is $O(k \log \frac{\log n \cdot \log |U|}{\delta_1} + \log \frac{\log n \log\log n \cdot \log |U|}{\delta_2} + k_g \log n)$. In most allocation schemes, we choose $k \leq \log\log n, k_g = O(\log\log n), \delta_1 = 1/\mathsf{poly}(n)$, and $\delta_2 = (\log n)^{-O(\log n)}$ such that the seed length is $O(\log n \log\log n)$.

Remark 4.2 *Notice that our parameters of $h_1 \circ h_2 \cdots \circ h_{k+1}$ are stronger than the parameters in [CRSW13]. In [CRSW13], the last function h_{k+1} is still a δ_1-biased space. But we are going to use $\delta_2 = (\delta_1)^{O(k)}$ in h_{k+1}, which provides almost $O(\log n)$-wise independence on $(\log n)^{O(\log n)}$ subsets of size $O(\log n)$ for our calculations.*

The main property of $h_1 \circ h_2 \cdots \circ h_k$: We fix g and discuss $h_1 \circ h_2 \cdots \circ h_k \circ h_{k+1}$ here. For each $i \in [k]$, it is natural to think about $h_1 \circ h_2 \cdots \circ h_i$ as a function from U to $[n^{1-\frac{1}{2^i}}]$, which is a hash map throw all balls into $n^{1-\frac{1}{2^i}}$ bins. Suppose the number of balls in one fixed bin is close to the expectation $n^{\frac{1}{2^i}} \cdot \frac{m}{n}$. Then h_{i+1} further allocates about $n^{\frac{1}{2^i}} \cdot \frac{m}{n}$ balls in this bin into $n^{\frac{1}{2^{i+1}}}$ bins after fixing $h_1 \circ h_2 \cdots \circ h_i$.

Celis et al. [CRSW13] showed that for every $i \in [k]$, the number of balls in every bin of $h_1 \circ h_2 \cdots \circ h_i$ is close to its expectation $n^{\frac{1}{2^i}} \cdot \frac{m}{n}$ in $\frac{1}{\mathsf{poly}(n)}$-biased spaces. We provide a proof here for completeness. We set $k = \log_2(\log n/3 \log\log n)$ and $\alpha = (\log n)^{-0.2}$ in this proof.

Lemma 4.3 *([CRSW13]) For any constant $c > 0$, there exists $\delta_1 = 1/\mathsf{poly}(n)$ such that given $m = O(n)$ balls, with probability at least $1 - n^{-c}$ over hash functions $h_1, \cdots, h_k$, for all $i \in [k]$, every bin in $[n^{1-\frac{1}{2^i}}]$ contains at most $(1+\alpha)^i n^{\frac{1}{2^i}} \cdot \frac{m}{n}$ balls under $h_1 \circ h_2 \cdots \circ h_i$.*

Proof. We apply induction on i. The base case $i = 0$ is true because there are at most m balls.

Suppose it is true for $i = l < d$. For a fixed bin $j \in [n^{1-\frac{1}{2^l}}]$, there are at most $(1+\alpha)^l n^{\frac{1}{2^l}} \cdot \frac{m}{n}$ balls. With out loss of generality, we assume there are exactly $s = (1+\alpha)^l n^{\frac{1}{2^l}} \cdot \frac{m}{n}$ balls from the induction hypothesis. Under the hash function h_{l+1}, we allocate these balls into $t = n^{\frac{1}{2^{l+1}}}$ bins.

We fix one bin in h_{l+1} and prove that this bin receives at most

$$(1+\alpha)s/t = (1+\alpha) \cdot (1+\alpha)^l n^{\frac{1}{2^l}} / n^{\frac{1}{2^{l+1}}} \cdot \frac{m}{n} = (1+\alpha)^{l+1} n^{\frac{1}{2^{l+1}}} \cdot \frac{m}{n}$$

balls with probability $\leq 2n^{-c-2}$ in a $\log^3 n$-wise δ_1-biased space for $\delta_1 = 1/\mathsf{poly}(n)$ with a sufficiently large polynomial. We use $X_i \in \{0,1\}$ to denote the ith ball is in the bin or not. Hence $\mathrm{E}[\sum_{i\in[s]} X_i] = s/t$.

For convenience, we use $Y_i = X_i - \mathrm{E}[X_i]$. Hence $Y_i = 1 - 1/t$ w.p. $1/t$, o.w. $Y_i = -1/t$. Notice that $\mathrm{E}[Y_i] = 0$ and $|\mathrm{E}[Y_i^l]| \leq 1/t$ for any $l \geq 2$.

We choose $b = 2^l \cdot \beta = O(\log n)$ for a large even number β and compute the bth moment of $\sum_{i\in[s]} Y_i$ as follows.

$$\begin{aligned}
\Pr[|\sum_{i\in[s]} X_i| > (1+\alpha)s/t] &\leq \mathrm{E}_{\delta_1\text{-biased}}[(\sum_i Y_i)^b]/(\alpha s/t)^b \\
&\leq \frac{\mathrm{E}_U[(\sum_i Y_i)^b] + \delta_1 \cdot s^{2b} t^b}{(\alpha s/t)^b} \\
&\leq \frac{\sum_{i_1,\cdots,i_b} \mathrm{E}[Y_{i_1} \cdot Y_{i_2} \cdots Y_{i_b}] + \delta_1 \cdot s^{3b}}{(\alpha s/t)^b} \\
&\leq \frac{\sum_{j=1}^{b/2} \binom{b-j-1}{j-1} \cdot \frac{b!}{2^{b/2}} \cdot s^j (1/t)^j + \delta_1 \cdot s^{3b}}{(\alpha s/t)^b} \\
&\leq \frac{2^{b/2} b! \cdot (s/t)^{b/2} + \delta_1 \cdot s^{3b}}{(\alpha s/t)^b}
\end{aligned}$$

Because $s/t \geq n^{\frac{1}{2^k}} \geq \log^3 n, b \leq \beta 2^k \leq \frac{\beta \log n}{3 \log\log n} \leq (s/t)^{1/3}$ and $\alpha = (\log n)^{-0.2} < (s/t)^{0.1}$, we simplify it to

$$\left(\frac{2b^2 \cdot s/t}{(\alpha s/t)^2}\right)^{b/2} + \delta_1 \cdot s^{3b} \leq \left(\frac{2(s/t)^{2/3} \cdot s/t}{(s/t)^{1.8}}\right)^{b/2} + \delta_1 \cdot s^{3b}$$
$$= (s/t)^{(-0.1)\cdot b/2} + \delta_1 \cdot s^{3b} \leq (n^{\frac{2}{2^{l+1}}})^{-0.1\cdot\beta 2^l} + \delta_1 (n^{\frac{3}{2^l}})^{3\beta\cdot 2^l} = n^{-c-2} + \delta_1 n^{9\beta} \leq 2n^{-c-2}.$$

Finally we choose $\beta = 40(c+2) = O(1)$ and $\delta_1 = n^{-9\beta-c-2}$ to finish the proof. □

Corollary 4.4 *For any constant $c > 0$, there exists $\delta_1 = 1/\mathsf{poly}(n)$ such that given $m = n$ balls, with probability at least $1 - n^{-c}$ over hash functions $h_1, \cdots, h_k$, for any bin $j \in [n^{1-\frac{1}{2^k}}] = [n/\log^3 n]$, it contains at most $1.01 \cdot \log^3 n \cdot \frac{m}{n}$ balls.*

We discuss about the last function h_{k+1} now. We always choose $\delta_2 = (\log n)^{-C\log n}$ for a constant C which depends on different applications. A toy example is to choose $\delta_2 = (\log n)^{-10\log n}$. Notice that under $h_1 \circ \cdots \circ h_k$, there are at most $(1+\alpha)^k n^{\frac{1}{2^k}} \leq 1.01 \cdot \log^3 n$ balls in each bin. Hence for any $\leq 1.01 \log^3 n$ balls in one bin, h_{k+1} is close to a $\log n$-wise independent distribution on the balls in a fixed bin. The distance between h_{k+1} and $\log n$-wise independence over all $[n^{1-\frac{1}{2^k}}]$ bins is at most

$$n^{1-\frac{1}{2^k}} \cdot \binom{1.01\log^3 n}{\log n} \cdot \delta_2 \cdot (\log^3 n)^{\log n/2} < \log n^{-2\log n}.$$

Remark 4.5 (Evaluation time) *Our hash function has an evaluation time $O((\log\log n)^4)$ in the RAM model from [MRRR14]. The reason is as follows.*

g can be evaluated by a degree $O(\log\log n)$ polynomial in the Galois field of size $\mathsf{poly}(n)$, which takes $O(\log\log n)$ time in the RAM model. The first k hash functions $h_1, \cdots, h_k$ still use $1/\mathsf{poly}(n)$-biased space, which have total evaluation time $O(k \cdot \log\log n) = O(\log\log n)^2$ in the RAM model.

The last function h_{k+1} in the RAM model is a $\log n$-wise $n^{-O(\log\log n)}$ biased space from U to $[\log^3 n]$, which needs $O(\log\log n)$ words in the RAM mode. Thus the evaluation time becomes $O(\log\log n)$ times the cost of a quadratic operation in the Galois field of size $n^{O(\log\log n)}$, which is $O((\log\log n)^4)$ from [MRRR14].

5 The *Uniform Greedy* scheme

We prove Theorem 1.1 in this section, i.e., given $m = O(n)$ balls in the *Uniform Greedy* scheme, the hash function of seed length $O(\log n \log\log n)$ in Section 4 achieves a maximum load of $\log_d \log n + O(1)$ with high probability.

Theorem 5.1 *For any $m = O(n)$, any constants $c > 0$ and $d \geq 2$, there exists a hash function of seed length $O(\log n \log\log n)$ such that for any m balls in U, with probability at least $1 - n^{-c}$, the max-load of the* Uniform Greedy *scheme with d independent choices from h is $\log_d \log n + O(c) + O(1)$.*

Proof. We bound the probability of a witness tree of height $\log_d \log n + \log_d(2+c) + 1 + 3c$ whose leaves have height at least $4d \cdot \frac{m}{n} + 1$ in d independent choices of our hash function h. We specify the parameter of h as follows: $k_g = 4c(\log_d \log n + \log_d(2+c) + 1 + 3c), k = \log_2(\log n/3\log\log n), \delta_2 = \log n^{-C\log n}$ for a large constant C, and $\delta_1 = 1/\mathsf{poly}(n)$ such that Lemma 4.3 and Corollary 4.4 holds with probability

at most n^{-c-1}. Let $h^{(1)}, \cdots, h^{(d)}$ denote the d independent hash functions from Setion 4 with the above parameters, where each

$$h^{(j)}(x) = \left(h_1^{(j)}(x) \circ h_2^{(j)}(x) \cdots \circ h_k^{(j)}(x) \circ h_{k+1}^{(j)}(x)\right) \oplus g^{(j)}(x).$$

We abuse the notation of g to denote $\{g^{(1)}, g^{(2)}, \cdots, g^{(d)}\}$ in the d choices and h_i to denote the group of hash functions $\{h_i^{(1)}, \cdots, h_i^{(d)}\}$ in this proof. We separate witness trees into two cases depending on the number of collisions.

Witness trees with at least $3c$ collisions: Given a witness tree T with at least $3c$ collisions, we consider the first $3c$ collisions $e_1, \cdots, e_{3c}$ in the BFS of T. Let T' be the subtree of T that only contains all vertices on edges $e_1, \cdots, e_{3c}$ and their ancestors in T. Therefore T survives under $h^{(1)}, \cdots, h^{(d)}$ only if T' survives under $h^{(1)}, \cdots, h^{(d)}$.

Observe that $|T'| \leq 3c \cdot 2 \cdot \text{height}(T)$. Because $k_g \geq |T'| + 3c - 1$ (the number of edges in T'), we bound the probability that T' survives in g is at most

$$(\frac{d}{n})^{|T'|+3c-1}.$$

At the same time, there are at most $m^{|T'|}$ choices of balls in T' and $(|T|^2)^{3c} = \mathsf{poly}(\log n)$ configurations of T'. Therefore we bound the probability that any witness with at least $3c$ collisions survives in our hash function by

$$(\frac{d}{n})^{|T'|+3c-1} \cdot m^{|T'|} \cdot |T|^{2\cdot 3c} \leq 0.5n^{-c}.$$

Witness tress with less than $3c$ collisions: We fix a configuration C of witness trees with height $\log_d \log n + \log_d(2+c)+1+3c$ but less than $3c$ collisions. Let $l = \lceil \log_d \log n + \log_d(2+c) \rceil$ with $d^l \in [(2+c)\log n, (2+c)d \log n]$ such that there are at least d^l leaves in this configuration. At the same time, $|C| \leq d^{l+2+3c}$ and the number of leaves is at least $\frac{|C|-3c}{2}$.

We consider all possible witness trees in this configuration. We extensively use the fact that after fixing $h_1 \circ \cdots \circ h_k$, at most $d(1.01 \log^3 n \cdot \frac{m}{n})$ elements in $h^{(1)}, \cdots, h^{(d)}$ are mapped to any one bin of $[n/\log^3 n]$ from Lemma 4.3.

We restate the condition (3) of the witness trees on the prefixes of $h^{(1)}, \cdots, h^{(d)}$ for every non-leaf ball u and its ith child v:

$$h_1^{(i)}(u) \circ \cdots \circ h_k^{(i)}(u) \in \left\{h_1^{(1)}(v) \circ \cdots \circ h_k^{(1)}(v), \cdots, h_1^{(d)}(v) \circ \cdots \circ h_k^{(d)}(v)\right\}. \tag{4}$$

Let $\mathcal{T}$ be the subset of witness trees in the configuration C whose edges satisfy the condition (4) in $h_1 \circ h_2 \cdots \circ h_k$, i.e., $\mathcal{T} = \{T | \text{configuration}(T) = C \text{ and } (u,v) \text{ satisfies (4) } \forall (u,v) \in T\}$. We claim

$$|\mathcal{T}| \leq m \cdot (d \cdot 1.01 \log^3 n \cdot \frac{m}{n})^{|C|-1}.$$

There are m choices of the root u in the witness tree. For the ith child v of the root u, we have to satisfy the condition (4) for (u,v). For a fixed bin $h_1^{(i)}(u) \circ \cdots \circ h_k^{(i)}(u)$, there are at most $1.01 \log^3 n \cdot \frac{m}{n}$ elements for each hash function $h^{(j)}$ that can be mapped to this bin from Lemma 4.3. Hence there are at most $d \cdot 1.01 \log^3 n \cdot \frac{m}{n}$ choices for each child of u. Then we repeat this arguments for all non-leaf nodes in the tree.

We first consider h_{k+1} as a uniform distribution from U to $[\log^3 n]$. We fix a witness tree T from $\mathcal{T}$. The probability that under the last part h_{k+1} of these hash functions $h^{(1)}, \cdots, h^{(d)}$, each edge (u, v) in $T \in \mathcal{T}$ satisfies $h^{(i)}(v) \in \{h^{(1)}(u), \cdots, h^{(d)}(u)\}$ is the same as the probability of

$$h_{k+1}^{(i)}(v) \in \left\{h_{k+1}^{(1)}(u), \cdots, h_{k+1}^{(d)}(u)\right\}, \tag{5}$$

which is at most $\frac{d}{\log^3 n}$.

For each leaf v in T, we claim the probability that its height is at least $4d \cdot \frac{m}{n} + 1$ is at most $2^{-d^2} \cdot (\frac{n}{m})^d$. Given a choice $i \in [d]$ of leaf v, we fix the bin to be $h^{(i)}(v)$. Then we bound the probability that there are at least $b = 4d \cdot \frac{m}{n}$ balls $w_1, \cdots, w_b$ in this bin excluding all balls in the tree by

$$\sum_{w_1 : w_1 < v, w_1 \notin T} \sum_{w_2 : w_1 < w_2 < v, w_2 \notin T} \cdots \sum_{w_b : w_{b-1} < w_b < v, w_b \notin T} \Pr[h^{(i)}(v) = h^{(j_1)}(w_1) = \cdots = h^{(j_b)}(w_b)]$$
$$\leq \frac{\binom{1.01d \cdot \log^3 n \cdot \frac{m}{n}}{b}}{(\log^3 n)^b} \leq \frac{(1.01d \cdot \frac{m}{n})^b}{b!} \leq (\frac{3}{4})^b.$$

Over d leaves, this probability is at most $(\frac{3}{4})^{b \cdot d} \leq 2^{-3d^2} \cdot (\frac{n}{m})^{2d}$.

Because $w_1, \cdots, w_b$ are not in the tree for every leaf, they are disjoint with the events in (5) which are over all edges in the tree. Hence we can times the above two probability together for each witness tree and apply the union bound on all witness trees in $\mathcal{T}$. Therefore the probability (in the uniform distribution) that there is one witness tree of height l whose leaves have height at least $4d \cdot \frac{m}{n} + 1$ is at most

$$|\mathcal{T}| \cdot (\frac{d}{\log^3 n})^{|C|-1} \cdot (2^{-2d^2} \cdot (\frac{n}{m})^{2d})^{\frac{|C|}{3}} \leq n \left(1.01d \cdot \log^3 n \cdot \frac{m}{n} \cdot \frac{d}{\log^3 n}\right)^{|C|} \cdot (2^{-2d^2} \cdot (\frac{n}{m})^{2d})^{|C|/3}$$
$$\leq n \cdot \left(2d^2 \cdot \frac{m}{n}\right)^{|C|} \cdot \left(2^{-2d^2} \cdot (\frac{n}{m})^{2d}\right)^{|C|/3} \leq n \cdot 2^{-d^l} \leq n^{-c-1}.$$

Now we choose $\delta_2 = (2d^2 \log n)^{-20(2+c)b \cdot d^{4c} \log n}$. Therefore in δ_2-biased spaces, we bound the probability that there is one witness tree of height $l + 3c$ whose leaves have height at least $4d \cdot \frac{m}{n} + 1$ by

$$(n^{-c-1} + |\mathcal{T}| \cdot \delta_2 \cdot (\log^3 n)^{|C| + b \cdot d^{l+1}}) < n^{-c-1}.$$

Then we apply a union bound on all possible figurations with at most $3c$ collisions:

$$(d^{l+2+3c})^{|3c|} \cdot n^{-c-1} \leq 0.5 n^{-c}.$$

From all discussion above, with probability at least $1 - n^{-c}$, there is no ball of height more than $\log_d \log n + 4c + 4d \cdot \frac{m}{n} + 2$. □

6 The *Always-Go-Left* Scehme

We show that given $m = O(n)$ balls in the *Always-Go-Left* scheme [VÖ03] with d choices, the hash function in Section 4 also achieves a max-load of $\frac{\log \log n}{d \log \phi_d} + O(1)$, where $\phi_d > 1$ is the constant satisfying $\phi_d^d = 1 + \phi_d + \cdots + \phi_d^{d-1}$. We define the *Always-Go-Left* scheme [VÖ03] as follows:

Definition 6.1 (*Always-Go-Left* with d choices) *Our algorithm partition the bins into d groups $G_1, \cdots, G_d$ of the same size n/d. Let $h^{(1)}, \cdots, h^{(d)}$ be d functions from U to $G_1, \cdots, G_d$ separately. For each ball i, the algorithm consider d bins $\{h^{(1)}(i) \in G_1, \cdots, h^{(d)}(i) \in G_d\}$ and chooses the bin with the least number of balls. If there are several bins with the least number of balls, our algorithm always choose the bin with the smallest group number.*

We define the asymmetric witness tree for the *Always-Go-Left* mechanism such that a ball of height $l + C + 1$ in any bin of $[n]$ indicates that there is an asymmetric witness tree of "height" $l + 1$ whose leaves have height at least C. There are two differences between asymmetric witness trees in this section and symmetric witness trees in the last section. In this section, the "height" of an asymmetric witness tree is the length of the shortest path from the root to the leaves. At the same time, there are d different types of asymmetric witness trees of the same "height" because the scheme is asymmetric in the d groups.

Definition 6.2 (Asymmetric Witness tree) *The asymmetric witness tree T of "height" l in group G_i is a d-ary tree. The root has d children where the subtree of the jth child is an asymmetric witness tree in group G_j of "height" $(l - 1_{j \geq i})$.*

Given d functions $h^{(1)}, \cdots, h^{(d)}$ from U to $G_1, \cdots, G_d$ separately, a ball b with height more than $l + 1$ in a bin belongs to group G_i indicates an asymmetric witness tree T of "height" l in G_i. Each node of T corresponds to a ball, and the root of T corresponds to the ball b. A ball u in T has a ball v as its jth child iff when we insert the ball u in the Always-Go-Left *mechanism, v is the top ball in the bin $h^{(j)}(u)$. Hence $v < u$ and $h^{(j)}(u) = h^{(j)}(v)$ when the jth child of u is v.*

For an asymmetric witness tree T of "height" l in group G_i, We use the "height" l and the group index $i \in [d]$ to determine its size. Let $f(l, i)$ be the size of a *full* asymmetric witness tree of "height" l in group G_i. From the definition, we have

$$f(l, i) = \sum_{j=1}^{i-1} f(l, j) + \sum_{j=i+1}^{d} f(l - 1, j).$$

Let $g\big((l-1) \cdot d + i\big) = f(l, i)$ such that

$$g(n) = g(n - 1) + g(n - 2) + \cdots + g(n - d).$$

We know there exist $c_0 > 1/2$ and $\phi_d > 1$ satisfying

$$\phi_d^2 = 1 + \phi_d + \cdots + \phi_d^{d-1} \text{ such that } g(n) \geq c_0 \cdot \phi_d^n.$$

Hence there exists a constant $c_0 > 1/2$ such that

$$f(l, i) = g\big((l - 1) \cdot d + i\big) \geq c_0 \cdot \phi_d^{(l-1)d+i}.$$

Vöcking in [VÖ3] showed that in a perfectly random hash function, the maximum load is $\frac{\log \log n}{d \log \phi_d} + O(1)$ with high probability given any n balls. We outline Vöcking's argument for *distinct balls* here: let b be a ball of height $l + 4$ for $l = \frac{\log \log n + \log(1+c)}{d \log \phi_d} + 1$. Without loss of generality, we assume that b is in the first group G_1. By the definition of the asymmetric witness tree, there exists a tree T in G_1 with root b and height l whose leaves have height at least 4. For each ball u and its ith ball v, $h^{(i)}(u) = h^{(i)}(v)$. We apply a union bound on all possible witness trees of height l in group G_1 to bound the probability by

$$n^{f(l,1)} \cdot (\frac{d}{n})^{f(l,1)-1} \cdot (\frac{1}{3^d})^{\text{number of leaves in } f(l,1)},$$

which is less than n^{-c} given $f(l,1) = \Theta(\phi_d^{(l-1)d+1}) = \Theta\big((1+c)\log n\big)$

We prove our derandomization of Vöcking's argument here.

Theorem 6.3 *For any $m = O(n)$, any constants $c > 0$ and $d \geq 2$,, there exist a constant $\phi_d \in (1.61, 2)$ and a hash function of seed length $O(\log n \log\log n)$ such that for any m balls in U, with probability at least $1 - n^{-c}$, the max-load of the* Always-Go-Left *mechanism with d independent choices from h is $\frac{\log\log n}{d\log\phi_d} + O(1)$.*

Proof. Let l be the smallest integer such that $c_0\phi_d^{ld} \geq (1+c)\log n$ and $b = 4d\frac{m}{n} + 1$. We bound the probability of a witness tree of height l whose leaves have height more than b in our hash function h during the *Always-Go-Left* scheme. Notice that there is a ball of height $l+b+3c+1$ in any bin of $G_2, G_3, \cdots, G_d$ indicates that there is a ball of the same height in G_1.

We choose our parameters of h as follows: $k_g = 2d\cdot(l+b+1+3c) = O(\log\log n)$, $k = \log_2(\log n/3\log\log n)$, $\delta_1 = 1/\mathsf{poly}(n)$ such that Lemma 4.3 happens with probability at most n^{-c-1}, and the bias $\delta_2 = \log n^{-O(\log n)}$ of h_{k+1} later. We set h_{k+1} to be a hash function from U to $[\log^3 n/d]$ and g to be a function from U to $[n/d]$ such that

$$h^{(j)} = \big(h_1^{(j)} \circ h_1^{(j)} \circ \cdots h_k^{(j)} \circ h_{k+1}^{(j)}\big) \oplus g^{(j)}$$

is a map from U to G_j of $[n/d]$ bins for each $j \in d$.

We use $h^{(1)}, \cdots, h^{(d)}$ to denote d independent hash functions with the above parameters. We abuse the notation of h_i to denote the group of hash functions $\{h_i^{(1)}, \cdots, h_i^{(d)}\}$ in this proof. We assume Lemma 4.3 and follow the same argument in Theorem 5.1. We bound the probability of witness trees from 2 cases depending on the number of collisions.

Witness trees with at least $3c$ collisions: Given a witness tree T with at least $3c$ collisions, we consider the first $3c$ collisions $e_1, \cdots, e_{3c}$ in the BFS of T. Let T' be the subtree of T that only contains all vertices in $e_1, \cdots, e_{3c}$ and their ancestors in T. Therefore T survives under $h^{(1)}, \cdots, h^{(d)}$ only if T' survives under $h^{(1)}, \cdots, h^{(d)}$.

Observe that $|T'| \leq 3c \cdot 2 \cdot \big(d \cdot \text{“height”}(T)\big)$. Because $k_g \geq |T'| + 3c$, we bound the probability that T' survives in g by

$$(\frac{d}{n})^{|T'|+3c-1}.$$

At the same time, there are at most $m^{|T'|}$ choices of balls in T' and $(|T|^2)^{3c} = \mathsf{poly}(\log n)$ configurations of T'. Hence we bound the probability of any witness with at least $3c$ collisions surviving by

$$(\frac{d}{n})^{|T'|+3c-1} \cdot m^{|T'|} \cdot |T|^{2\cdot 3c} \leq n^{-c+1}.$$

Witness tree with less than $3c$ collisions: We fix a configuration C of witness tree in group G_1 with height $l+1+3c$ and less than $3c$ collisions and consider all possible witness tree in this configuration C. Thus $|C| \in [f(l+1,1), f(l+1+3c,1)]$.

Let $\mathcal{T}$ be the subset of possible asymmetric witness tree in C. For any $T \in \mathcal{T}$, each edge (u,v) has to satisfy $h^{(i)}(u) = h^{(i)}(v)$ in the *Always-Go-Left* scheme when v is the ith child of u, which bounds

$$|\mathcal{T}| \leq n \cdot (1.01\log^3 n \cdot \frac{m}{n})^{|C|-1}$$

under $h_1 \circ h_2 \cdots \circ h_k$ based on Lemma 4.3.

We first consider h_{k+1} as a uniform distribution from U to $[\log^3 n/d]$ then move to the δ_2-biased space. For each asymmetric witness tree, every edge (u, v) maps to the same bin w.p. $d/\log^3 n$ in h_{k+1}.

For each leaf, its height is at least b if each of its choice has height at least $b-1$, which happens with probability at most

$$\left(\frac{\binom{1.01\cdot\log^3 n\cdot\frac{m}{n}}{b-1}}{(\log^3 n/d)^{b-1}}\right)^d \le \left(\frac{(1.01d\cdot\frac{m}{n})^{b-1}}{(b-1)!}\right)^d \le 2^{-d^2}\cdot(\frac{n}{m})^d.$$

Because these two types of events are on disjoint subsets of balls, an asymmetric witness tree of "height" at least $l+1$ in group G_1 whose leaves have height at least b happens with probability at most

$$\begin{aligned}|\mathcal{T}|\cdot\left(\frac{d}{\log^3 n}\right)^{|C|-1}\cdot\left(2^{-d^2}\cdot(\frac{n}{m})^d\right)^{\frac{(d-1)(|C|-3c)}{d}} \le& n\cdot\left(1.01d\cdot\frac{m}{n}\right)^{|C|}\cdot\left(2^{-d^2}\cdot(\frac{n}{m})^d\right)^{|C|/3}\\ \le& n\cdot 2^{-f(l+1,1)}\le n^{-c-1}.\end{aligned}$$

We choose $\delta_2=(2^d\log n)^{-10(1+c)\cdot b\cdot 2^{3cd}\log n}$ such that in δ_2-biased spaces, this happens with probability at most $n^{-c-1}+|\mathcal{T}|\cdot\delta_2\cdot(\log^3/d)^{bd\cdot f(l+3c+1,1)}\le 2n^{-c-1}$. At the same time, the number of possible configurations is at most $(f(l+3c+1,1)^2)^{3c}\le 0.1n$.

From all discussion above, with probability at most n^{-c}, there exists a ball in the *Always-Go-Left* mechanism with height at least $l+b+1=\frac{\log n\log n}{d\log\phi_d}+O(1)$. □

7 Heavy load

We consider the derandomization of the 1-choice scheme when we have $m=n\mathsf{poly}(\log n)$ balls and n bins. From the Chernoff bound, w.h.p, the max-load among n bins is $\frac{m}{n}\left(1+O(\sqrt{\log n}\cdot\sqrt{\frac{n}{m}})\right)$ when we throw $m>n\log n$ balls into n bins independently at random. We modify the hash function from [CRSW13] with proper parameters for $m=\mathsf{poly}(\log n)\cdot n$ balls and prove the max-load is still $\frac{m}{n}\left(1+O(\sqrt{\log n}\cdot\sqrt{\frac{n}{m}})\right)$. We assume $m=\log^a n\cdot n$ for a constant $a\ge 1$ in the rest of this section.

Theorem 7.1 *For any constant $c>0$ and $a\ge 1$, there exist a constant C and a hash function from U to $[n]$ generated by $O(\log n\log\log n)$ random bits such that for any $m=\log^a n\cdot n$ balls, with probability at least $1-n^{-c}$, the max-load of the n bins in the 1-choice scheme with the hash function h is at most $\frac{m}{n}\left(1+C\cdot\sqrt{\log n}\cdot\sqrt{\frac{n}{m}}\right)$.*

We change our hash function with different parameters. We choose $k=\log\frac{\log n}{(2a)\log\log n}$, h_i to denote a hash function from U to $[n^{2^{-i}}]$ for $i\in[k]$, and h_{k+1} to denote a hash function from U to $[n^{2^{-k}}]=[\log^{2a} n]$ such that $h_1\circ h_2\cdots\circ h_k\circ h_{k+1}$ constitute a hash function from U to $[n]$. We set $\alpha=4(c+2)\sqrt{\log n}\sqrt{\frac{n}{m}}$. For convenience, we still think $h_1\circ h_2\cdots\circ h_i$ as a hash function maps to $n^{1-2^{-i}}$ bins for any $i\le k$. In this section, we use δ_1-biased spaces on $h_1,\cdots,h_k$ and a δ_2-biased space on h_{k+1} for ease of exposition.

Claim 7.2 *For any constant $c>0$, there exists $\delta_1=1/\mathsf{poly}(n)$ such that given $m=\log^a n\cdot n$ balls, with probability $1-n^{-c-1}$, for any $i\in[k]$ and any bin $b\in[n^{1-2^{-i}}]$, there are less than $\prod_{j\le i}(1+\frac{\alpha}{(k+2-i)^2})\cdot\frac{m}{n}\cdot n^{2^{-i}}$ balls in this bin.*

Proof. We still use induction on i. The base case is $i = 0$. Because there are at most m balls, the hypothesis is true.

Suppose it is true for $i = l$. Now we fix a bin and assume there are $s = \prod_{j\le l}(1 + \frac{\alpha}{(k+2-i)^2}) \cdot \frac{m}{n} n^{2^{-l}} \le (1+\alpha)\frac{m}{n} n^{2^{-l}}$ balls in this bin from the induction hypothesis. h_{l+1} maps these s balls to $t = n^{2^{-(l+1)}}$ bins. We will prove that with high probability, every bin in these t bins of h_{l+1} contains at most $(1 + \frac{\alpha}{(k+1-l)^2})s/t$ balls.

We use $X_i \in \{0, 1\}$ to denote whether ball i is in one fixed bin of $[t]$ or not. Hence $\Pr[X_i = 1] = 1/t$. Let $Y_i = X_i - \mathrm{E}[X_i]$. Therefore $\mathrm{E}[Y_i] = 0$ and $\mathrm{E}[|Y_i|^l] \le 1/t$ for any $l \ge 2$. Let $b = \beta 2^l$ for a large constant β later.

$$\begin{aligned}
\Pr_{D_{\delta_1}}[\sum_i X_i > (1 + \frac{\alpha}{(k+1-l)^2})s/t] &\le \frac{\mathrm{E}_{D_{\delta_1}}[(\sum_i Y_i)^b]}{(\frac{\alpha}{(k+1-l)^2}s/t)^b} \\
&\le \frac{\sum_{i_1,\cdots,i_b} \mathrm{E}_U[Y_{i_1}\cdots Y_{i_b}] + \delta_1 s^{2b}}{(\frac{\alpha}{(k+1-l)^2}s/t)^b} \\
&\le \frac{2^b b!(s/t)^{b/2} + \delta_1 s^{2b}}{(\frac{\alpha}{(k+1-l)^2}s/t)^b} \\
&\le \left(\frac{2b(s/t)}{(\frac{\alpha}{(k+1-l)^2}s/t)^2}\right)^{b/2} + \delta_1 \cdot s^{2b}
\end{aligned}$$

We use these bounds $k = \log \frac{\log n}{(2l)\log\log n} < \log\log n$, $b < \beta 2^k < \frac{\beta \log n}{(2l)\log\log n}$ and $n^{2^{-l-1}} \ge n^{2^k} \ge \log^{2l} n \ge (m/n)^2$ to simplify the above bound by

$$\begin{aligned}
&\left(\frac{2\log n}{\frac{\alpha^2}{(\log\log n)^4}\cdot s/t}\right)^{b/2} + \delta_1 s^{2b} \\
\le &\left(\frac{2\log^2 n}{(\log n \cdot \frac{n}{m}) \cdot (\frac{m}{n} n^{2^{-l-1}})}\right)^{b/2} + \delta_1 s^{2b} \\
\le &\left(\frac{1}{n^{0.5\cdot 2^{-l-1}}}\right)^{b/2} + \delta_1 s^{2b} \\
\le &n^{-0.5\cdot 2^{-l-1}\cdot \beta 2^l/2} + \delta_1 \left(\frac{2m}{n} n^{2^{-l}}\right)^{2\beta 2^l} \le n^{-\beta/8} + \delta_1 \cdot n^{6\beta}.
\end{aligned}$$

Hence we choose the two parameters $\beta > 8(c+2)$ and $\delta_1 = n^{-6\beta-c-2}$ such that the above probability is bounded by $2n^{-c-2}$. Finnally, we apply the union bound on i and all bins. □

Proof of Theorem 7.1. We first apply Claim 7.2 to $h_1, \cdots, h_k$.

In h_{k+1}, we first consider it as a $b = 16(c+2)^2 \log n = O(\log n)$-wise independent distribution that maps $s < \prod_{j\le k}(1 + \frac{\alpha}{(k+2-i)^2}) \cdot \frac{m}{n} n^{2^{-k}}$ balls to $t = n^{2^{-k}}$ bins. From Lemma 3.4 and Theorem 5 (I) in [SSS95], we bound the probability that one bin receives more than $(1+\alpha)s/t$ by $e^{\alpha^2 \cdot \mathrm{E}[s/t]/3} \le n^{-c-2}$ given $b \ge \alpha^2 \mathrm{E}[s/t]$.

Then we choose $\delta_2 = (\log n)^{-b\cdot 5a} = (\log n)^{-O(\log n)}$ such that any δ_2-biased space from $[2\frac{m}{n}\log^{2a} n]$ to $[\log^{2a} n]$ is $\delta_2 \cdot \binom{2\frac{m}{n}\log^{2a} n}{\leq b} \cdot (\log^{2a} n)^b < n^{-c-2}$-close to a b-wise independent distribution. Hence in h_{k+1}, with probability at most $2 \cdot n^{-c-2}$, there is one bin that receives more than $(1+\alpha)s/t$ balls. Overall, the number of balls in any bin of $[n]$ is at most

$$\prod_{i\leq k}(1+\frac{\alpha}{(k+2-i)^2})(1+\alpha)\frac{m}{n} \leq (1+\sum_{i\leq k+1}\frac{\alpha}{(k+2-i)^2})\frac{m}{n} \leq (1+2\alpha)\frac{m}{n}.$$

□

Acknowledgement

The author is grateful to David Zuckerman for his constant support and encouragement, as well as for many fruitful discussions.

References

[ABKU99] Yossi Azar, Andrei Z. Broder, Anna R. Karlin, and Eli Upfal. Balanced allocations. *SIAM J. Comput.*, 29(1):180–200, September 1999. 0, 1, 2, 3

[ADM+99] Noga Alon, Martin Dietzfelbinger, Peter Bro Miltersen, Erez Petrank, and Gábor Tardos. Linear hash functions. *J. ACM*, 46(5):667–683, September 1999. 3

[AGHP90] N. Alon, O. Goldreich, J. Hastad, and R. Peralta. Simple construction of almost k-wise independent random variables. In *Proceedings of the 31st Annual Symposium on Foundations of Computer Science*. IEEE Computer Society, 1990. 8

[CRSW13] L. Elisa Celis, Omer Reingold, Gil Segev, and Udi Wieder. Balls and bins: Smaller hash families and faster evaluation. *SIAM J. Comput.*, 42(3):1030–1050, 2013. 0, 1, 2, 3, 4, 9, 10, 16

[DKRT16] Søren Dahlgaard, Mathias Bæk Tejs Knudsen, Eva Rotenberg, and Mikkel Thorup. The power of two choices with simple tabulation. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016*, pages 1631–1642, 2016. 2, 3

[DW03] Martin Dietzfelbinger and Philipp Woelfel. Almost random graphs with simple hash functions. In *Proceedings of the Thirty-fifth Annual ACM Symposium on Theory of Computing*, STOC '03, pages 629–638, 2003. 3

[GKM15] Parikshit Gopalan, Daniek Kane, and Raghu Meka. Pseudorandomness via the discrete fourier transform. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015*, pages 903–922, 2015. 2, 3

[Mit01] Michael Mitzenmacher. The power of two choices in randomized load balancing. *IEEE Trans. Parallel Distrib. Syst.*, 12(10):1094–1104, 2001. 3

[MRRR14] Raghu Meka, Omer Reingold, Guy N. Rothblum, and Ron D. Rothblum. Fast pseudorandomness for independence and load balancing - (extended abstract). In *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Proceedings*, pages 859–870, 2014. 1, 3, 7, 11

[MRS00] Michael Mitzenmacher, Andrea W. Richa, and Ramesh Sitaraman. The power of two random choices: A survey of techniques and results. In *in Handbook of Randomized Computing*, pages 255–312, 2000. 3

[NN90] J. Naor and M. Naor. Small-bias probability spaces: Efficient constructions and applications. In *Proceedings of the Twenty-second Annual ACM Symposium on Theory of Computing*, STOC '90, pages 213–223. ACM, 1990. 8

[PT12] Mihai Pătraşcu and Mikkel Thorup. The power of simple tabulation hashing. *J. ACM*, 59(3):14:1–14:50, June 2012. 3

[RRW14] Omer Reingold, Ron D. Rothblum, and Udi Wieder. Pseudorandom graphs in data structures. In *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Proceedings*, pages 943–954, 2014. 1, 2, 3

[SSS95] Jeanette P. Schmidt, Alan Siegel, and Aravind Srinivasan. Chernoff-hoeffding bounds for applications with limited independence. *SIAM J. Discret. Math.*, 8(2):223–250, May 1995. 3, 8, 17

[VÖ3] Berthold Vöcking. How asymmetry helps load balancing. *J. ACM*, 50(4):568–589, July 2003. 0, 1, 2, 3, 5, 8, 13, 14

[Vaz86] Umesh Vazirani. Randomness, adversaries and computation. In *Ph.D. Thesis, EECS, UC Berkeley*, pages 458–463. 1986. 4, 8

[Woe06] Philipp Woelfel. Asymmetric balanced allocation with simple hash functions. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithm*, SODA '06, pages 424–433, 2006. 2, 3